

REMARKS

Claims 1-28, 53 and 54 remain in this application.

In response to the restriction requirement, Claims 29-52 and 55-76 have been canceled.

Antecedent support for the amendments to the claims is found in Claim 16 and page 20, lines 3-5 and 15-24.

The Examiner has rejected Claims 1-6, 16-19, 25 and 26 as being anticipated by Jiang. Applicants respectfully traverse this rejection.

Jiang teaches that the data movers are data mover computers. Each of the data mover computers has a respective high-speed data link to a respective port of the cached disk array. The cached disc array is configured so that the filesystem is accessible only through the dataport connected to the first data mover and so that the filesystem is accessible only through the dataport connected to the second data mover. Each of the data movers maintains a directory of the data mover ownership of all of the files in the first and second filesystems. In other words, each of the data movers maintains a copy of the filesystem configuration information in order to recognize which data mover in the system has exclusive access to a specified read/write file.

Each of the data movers may receive file access requests from at least one network client. The clients communicate with the data movers using a connection oriented NFS protocol. Whenever the data mover receives a file access request from the client, it checks the configuration directory to determine whether or not the file specified by the request is in a filesystem owned by the data mover. If so, then the data mover places a lock on the specified file, accesses the file in the filesystem, and streams read/write data between the client and the filesystem. If the file specified by the request is not a filesystem owned by the data mover, the data mover forwards the request to the data mover that owns the filesystem to be accessed. See column 1, line 63-column 2, line 49.

A further modification of the above architecture is achieved by including a data bypass path between the first data mover in the second filesystem in order to bypass the second data mover, and a data bypass path between the second data mover and the first filesystem in order to bypass the first data mover. It is possible for each of the data movers to access data in each of the filesystems, but if a data mover does not have the file access information for the filesystem to be accessed, and the data mover to ask the owner for permission to access the filesystem, or else a data consistency problem arises. See column 8, lines 30-42.

As is clearly evident from the above description, a data mover computer is not a switching fabric, as found in Claim 1. In fact, Jiang teaches to have bypass paths to connect a data mover to a filesystem precisely so it does not need to act as a switch. The very teaching of bypass paths connecting data movers to filesystems teaches away from the data movers acting as

switches, since if they acted as switches, they would not need the bypass paths. In fact, the plain language that Jiang uses is "data mover computer" and not a switch or a switching fabric. The language that the author uses must be given deference, unless there is a clear contrary definition provided to the standard definition.

As explained above, each of the data movers maintains a directory of the data mover ownership of all the files in the first and second filesystems in order to recognize which data mover in the system has exclusive access to a specified file. Whenever the data mover receives a file access request from the client, it checks the configuration directory to determine whether or not the file specified by the request is in a filesystem owned by the data mover. If the file specified by the request is not a filesystem owned by the data mover, the data mover forwards the request to the data mover that owns the filesystem to be accessed, not to the filesystem.

As is evident from the above description, a request, as taught by Jiang, must search a configuration directory in the data mover in order to determine whether the desired filesystem is owned by the data mover. Accordingly, the request has no idea where the filesystem which it desires is located. There is no teaching or suggestion of routing the request from the network element to the disk element through the switching fabric based on the filesystem request ID. This limitation represents the fact that a switching fabric found in claim 1 of applicants is a true switching fabric, as opposed to a data mover computer which has the capability of guiding the request, as taught by Jiang, but only after the data mover itself

determines where the filesystem is located and whether the data mover has ownership of the filesystem that corresponds to the request.

Furthermore, to emphasize the fact that the switching fabric is different from a data mover, there is also present the limitation in Claim 1 that the switching fabric processes higher priority requests before lower priority requests. There is no teaching or suggestion anywhere about the requests having any priority let alone being processed based on their priority.

In addition, to emphasize the true switching nature present in Claim 1, there is also the limitation of "a remote procedure call mechanism which forms a unique connection between a network element and a disk element through the switch fabric at a certain priority through which requests and responses between the disk element and network elements flow". There is no teaching or suggestion of a unique connection being formed between the client and data mover and filesystem in Jiang. Jiang simply teaches that each step of the way is isolated from the next step in that the request transfers from the client to the data mover and then from the data mover to the filesystem unconnected or unrelated to the connection between the client and the data mover, and vice versa. This further manifests the true switching nature of applicants' claimed invention.

Accordingly, Claim 1 is patentable over Jiang. Claims 2-15 are dependent to parent Claim 1 and are patentable for the reasons Claim 1 is patentable.

Claim 16 has the limitation of "routing the requests to a switching fabric through unique connections to the network elements based on the filesystem request ID according to the mapping function to disk elements connected to the switching fabric with the switching fabric processing higher priority requests before lower priority requests". As explained above, Jiang does not teach or suggest these limitations and Claim 16 is patentable over Jiang.

Claims 17-24 are dependent to parent Claim 16 and are patentable for the reasons Claim 16 is patentable.

Claim 25 is patentable for the reasons Claim 1 is patentable.

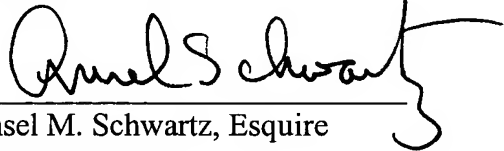
Claim 26 is patentable for reasons Claim 1 is patentable.

Claim 27 is patentable for the reasons Claim 1 is patentable.

Claim 52 is patentable for the reasons Claim 1 is patentable.

In view of the foregoing amendments and remarks, it is respectfully requested that the outstanding rejections and objections to this application be reconsidered and withdrawn, and Claims 1-28, 53 and 54, now in this application be allowed.

Respectfully submitted,

By 

Ansel M. Schwartz, Esquire

Reg. No. 30,587

One Sterling Plaza

201 N. Craig Street

Suite 304

Pittsburgh, PA 15213

(412) 621-9222

Attorney for Applicants